

**METHODS AND SYSTEMS FOR EDITING
A NETWORK OF INTERCONNECTED CONCEPTS**

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application is a continuation-in-part of United States Patent Application Serial No. 09/493,701 filed January 28, 2000 entitled "Meaning-Based Advertising and Relevance Determination," which is is a continuation-in-part of United States Patent No. 6,453,315 filed November 1, 1999 entitled "Meaning-Based Information Organization and Retrieval," which claims priority to United States Provisional Patent Application Serial No. 60/155,667 filed September 22, 1999, all of which are hereby incorporated in their entirety by this reference, and this application claims priority to United States Provisional Patent Application Serial No. 60/491,422 filed July 30, 2003 entitled "Systems and Methods of Organizing and Retrieving Information Based on Meaning," which is hereby incorporated in its entirety by this reference.

FIELD OF THE INVENTION

[0002] The invention generally relates to a network of interconnected concepts. More particularly, the invention relates to methods and systems for editing a network of interconnected concepts.

BACKGROUND OF THE INVENTION

[0003] Improving mechanisms for searching vast numbers of documents, such as those available via the World Wide Web (“web”) or on large enterprise computer networks, has increasingly been an area of focus and research. This is especially due to the continued growth in the number of computer users, services and applications offered for example on the web, and the amount of information being added to the web, enterprise computer networks and other databases.

[0004] A search engine or search engine program is a widely used mechanism for allowing users to search vast numbers of documents for information. Conventional search engines index the documents and the terms appearing in the documents. A term can be a string of characters or groups of strings of characters representing a one or more related words. Generally, a search engine provides a user interface that includes a query field. A user can enter one or more search terms describing the desired information in the query field. In response to a query, the search engine attempts to locate relevant documents by matching the search terms to the indexed terms appearing in the documents. Based on these matches, the search engine ranks, sorts and then returns for display the search results, usually as a list of relevant documents.

[0005] Typically, the concept or concepts represented by the search terms or the terms in the documents are not taken into account. A concept can be a normalized semantic representation that can be expressed with any number of terms. Methods of complex searching exist, but most of these methods of complex searching are term

based. For example, an advanced Boolean search or an advanced text search are both based on terms and not concepts. With a term search, documents are converted into a list of terms with relative rankings and the document's terms are compared with the search terms. Term searches can provide poor results when faced with terms having multiple meanings, the use of synonyms, and important terms that are not used frequently. Further, it is difficult to search for a concept using term-based searching methods. For example, if it is desired to search all documents that mention the concept "California cities", each city must be enumerated as a search term.

[0006] In order to provide a concept-based search, a search engine can utilize a network of concepts. While such networks of concepts exist, these conventional networks generally associate concepts in a hierarchical fashion and provide no weight to the associations. Further, such conventional networks of concepts are not easily editable.

SUMMARY

[0007] Embodiments of the present invention comprise methods and systems for editing a network of interconnected concepts. In one embodiment, a semantic engine implements a method comprising receiving a primary term representing a first concept to be added to a network of interrelated concepts, receiving at least one related term related to the primary term, receiving at least one relationship between the first concept and a second concept, receiving a strength value associated with the relationship, and adding the first concept to the network. In another embodiment, a

semantic engine implements a method comprising receiving a request to edit a first concept in a network of interrelated concepts, causing the display of the first concept, receiving a request to add at least one relationship between the first concept and a second concept, receiving a strength value associated with the relationship, and updating the first concept in the network. Additional aspects of the present invention are directed to computer systems and to computer-readable media having features relating to the foregoing aspects.

[0008] Further details and advantages of embodiments of the present invention are set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] These and other features, aspects, and advantages of the present invention are better understood when the following Detailed Description is read with reference to the accompanying drawings, wherein:

[0010] FIG. 1 illustrates a block diagram of a system in accordance with one embodiment of the present invention.

[0011] FIG. 2 illustrates an example of a portion of a relationship-based ontology in accordance with one embodiment of the present invention.

[0012] FIG. 3 illustrates an example of bond strength and semantic distance in a portion of a relationship-based ontology in accordance with one embodiment of the present invention.

[0013] FIG. 4 illustrates an exemplary concept view in accordance with one embodiment of the present invention.

[0014] FIG. 5 illustrates a flow diagram of a method in accordance with one embodiment of the present invention.

[0015] FIG. 6 illustrates a flow diagram of a method in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[0016] The present invention comprises methods and systems for editing a network of interconnected concepts. In one embodiment, a semantic engine allows for the creation and editing of concepts in a network of interconnected concepts or ontology. The concepts can be defined by a synset or a group of related terms, such as synonyms and acronyms, relationships, both hierarchical and lateral, with other concepts, and the strength of the relationships with other concepts. The semantic engine allows for the creation or editing of a concept by allowing a user to define the related terms, the relationships with other concepts, and the strength of the relationships.

[0017] Reference will now be made in detail to exemplary embodiments of the invention as illustrated in the text and accompanying drawings. The same reference numbers are used throughout the drawings and the following description to refer to the same or like parts.

[0018] Various systems in accordance with the present invention may be constructed. FIG. 1 is a diagram illustrating an exemplary system in which exemplary embodiments of the present invention may operate. The present invention may operate, and be embodied in, other systems as well.

[0019] The system 100 shown in FIG. 1 includes multiple client devices 102a-n, a server device 104 and a network 106. The network 106 shown can be the Internet. In other embodiments, other networks, such as an intranet may be used. Moreover, methods according to the present invention may operate in a single computer. The client devices 102a-n shown each include a computer-readable medium, such as a random access memory (RAM) 108 in the embodiment shown, coupled to a processor 110. The processor 110 executes a set of computer-executable program instructions stored in memory 108. Such processors may include a microprocessor, an ASIC, and state machines. Such processors include, or may be in communication with, media, for example computer-readable media, which stores instructions that, when executed by the processor, cause the processor to perform the steps described herein. Embodiments of computer-readable media include, but are not limited to, an electronic, optical, magnetic, or other storage or transmission device capable of providing a processor, such as the processor in communication with a touch-sensitive input device, with computer-readable instructions. Other examples of suitable media include, but are not limited to, a floppy disk, CD-ROM, magnetic disk, memory chip, ROM, RAM, an ASIC, a configured processor, all optical media, all magnetic tape or other magnetic media, or any other medium from which a computer

processor can read instructions. Also, various other forms of computer-readable media may transmit or carry instructions to a computer, including a router, private or public network, or other transmission device or channel, both wired and wireless. The instructions may comprise code from any computer-programming language, including, for example, C, C++, C#, Visual Basic, Java, Python, Perl, and JavaScript.

[0020] Client devices 102a-n may also include a number of external or internal devices such as a mouse, a CD-ROM, a keyboard, a display, or other input or output devices. Examples of client devices 102a-n are personal computers, digital assistants, personal digital assistants, cellular phones, mobile phones, smart phones, pagers, digital tablets, laptop computers, a processor-based device and similar types of systems and devices. In general, a client device 102a-n may be any type of processor-based platform connected to a network 106 and that interacts with one or more application programs. The client devices 102a-n shown include personal computers executing a user interface application program and/or a browser application program such as Internet Explorer™, version 6.0 from Microsoft Corporation, Netscape Navigator™, version 7.1 from Netscape Communications Corporation, and Safari™, version 1.0 from Apple Computer. Through the client devices 102a-n, users 112a-n can communicate over the network 106 with each other and with other systems and devices coupled to the network 106.

[0021] As shown in FIG. 1, server device 104 is also coupled to the network 106. The server device 104 shown includes a server executing a semantic engine program. The system 100 can also include multiple other server devices. Similar to

the client devices 102a-n, the server device 104 shown includes a processor 116 coupled to a computer readable memory 118. The server device 104 is depicted as a single computer system, but may be implemented as a network of computer processors. Examples of server device 104 are servers, mainframe computers, networked computers, a processor-based device and similar types of systems and devices. Client processors 110 and server processor 116 can be any of a number of well known computer processors, such as processors from Intel Corporation of Santa Clara, California and Motorola Corporation of Schaumburg, Illinois.

[0022] Memory 118 of the server device 104 contains a semantic engine application program, also known as a semantic engine 120. The semantic engine 120 manages and provides access to a network of interconnected concepts or ontology. In one embodiment, the semantic engine 120 also can index articles and concepts found in articles to enable the location and retrieval of information. Articles include, documents, for example, web pages of various formats, such as HTML, XML, XHTML, Portable Document Format (PDF) files, and word processor, database, and application program document files, audio, video, or any other information of any type whatsoever made available on a network (such as the Internet or intranet), a personal computer, or other computing or storage means. The embodiments described herein are described generally in relation to documents, but embodiments may operate on any type of article.

[0023] The semantic engine 120 shown includes a concept processor 130, which, in the embodiment shown, comprises computer code residing in the memory

118. The concept processor 130 can edit and create concepts in a network of interconnected concepts based on information received from a client device 102a input by a user 112a. The client processor 130 can have a user interface application that allows user 112a via the client device 102a to interact with the concept processor 130 to edit existing concepts and create new concepts in the ontology.

[0024] Server device 104 also provides access to other storage elements, such as a concept storage element, in the example shown a concept database 122. The concept database 120 can be used to store the ontology. The concept database can be located in memory 118 as shown or can be located external to server device 104. Data storage elements may include any one or combination of methods for storing data, including without limitation, arrays, hash tables, lists, and pairs. Other similar types of data storage devices can be accessed by the server device 104.

[0025] It should be noted that the present invention may comprise systems having different architecture than that which is shown in FIG. 1. For example, in some systems according to the present invention, the semantic engine 120 may not be on a single server device 104, but may be on multiple server devices. The system 100 shown in FIG. 1 is merely exemplary, and is used to explain the exemplary methods shown in FIGs. 5 and 6.

[0026] As explained above, the concept database 122 can contain an ontology that can describe a semantic space that can be employed to facilitate concept-based location and retrieval of information. The ontology consists of a plurality of concepts, which are connected to one another through many different types of

relationships, such as, for example, “kind of”, “has kind”, “part of”, “has part”, “member of”, “has member”, “substance of”, “has substance”, “product of”, and “has product”. For example, the concept of “table” is connected to the concept of “furniture” through a “kind of” connection. Thus, “table” is a “kind of” “furniture”. Similarly, “California” is a part of the “United States”, the “United States” has part “California”, the “United States” is a member of the “NATO”, and “NATO” has member “United States”.

[0027] The ontology according to one embodiment of the present invention also includes bind or lateral bond relationships between concepts. A lateral bond relationship describes one concept’s closeness to another in an average person’s common understanding. For example, “skier” and “skiing” are not closely related concepts in conventional ontologies. The former is a kind of “athlete”, ultimately a kind of “human being”; and thus would reside within the “entity” or “living thing” tree. The latter is a kind of “sport”, ultimately a kind of “activity”; it is in the “actions” tree. Though the subjects are closely related in everyday usage, they may be in widely separated locations within a conventional ontology. To remedy this, a bind or lateral bond can be made between the two concepts, to reflect their close proximity in semantic space (when an average person thinks of one concept, the person tends to think of the other).

[0028] The relationships between elements may take on many forms and can become quite complex, but for ease of illustration, a simple example is illustrated by FIG. 2. The boxed elements in FIG. 2 represent meanings within the ontology and

collectively, along with the relationship connections between meanings can be viewed as defining a semantic space. The relationship types “part of”, “kind of” and “lateral bond” are represented by differing line types in FIG. 2, a legend for which is drawn thereon.

[0029] The example in FIG. 2 concerns skiing. Starting with the branch for “sport”, “skiing” is defined in the ontology 200 as a kind of “sport”. The word “ski” typically, in its noun form, can be thought of as related to “skiing” in that it is a “part of” “skiing” as shown in FIG. 2. “Slalom skiing” is a type of skiing and hence a “kind of” connection is shown between it and “skiing”. “Bindings” are a structural attachment on a ski, and hence it is assigned a “part of” connection with “ski”. The example of a specific brand of ski, “K2 ski,” is given to show how it is in a “kind of” connection with “ski”.

[0030] “K2 ski” is also assigned a lateral bond showing a conceptual commonness with the manufacturer of the ski “K2” which lies in the “company” branch. The lateral bond can be, for example, “product of”. The company branch has a child “athletic equipment company” as a “kind of” “company.” “Athletic equipment company” has as its child in turn the “K2” company.

[0031] Considering “ski” once again, “ski” is also a child of the “equipment” branch which has “athletic equipment” as a “kind of” “equipment” and ski as a “kind of” “athletic equipment”. “Surfboard” is related to “ski” in that it too is a “kind of” “athletic equipment”. Target documents or nodes within a subject directory may be “placed” or “located” by human intervention or automatically into the semantic space

as defined by ontology 200. A document containing information concerning the sale of skis or skiing destinations would fall somewhere within the defined semantic space based upon its focus of content.

[0032] The ontology according to one embodiment of the present invention also includes bond strength between concepts. A value can be assigned to a connection from one concept to another that signifies how strongly the second meaning relates to the first. These connections can be dependant on the direction of the bond, so that, for example, “skier” might imply a strong connection to “skiing”, whereas “skiing need not imply “skier” to the same degree.

[0033] FIG. 3 illustrates an example of bond strengths and semantic distances between the concepts illustrated in FIG. 2. FIG. 3 illustrates how distance and closeness of meaning between concepts can be quantified within the semantic space. Distances are shown between the element “ski” and all other elements within the semantic space. Using three classes of bond strengths the degree of closeness between concepts may be discovered. A “strong relationship” exists between “ski” and “skiing” as does between “ski” and “athletic equipment.” Between “skiing” and “sport” there is a weaker than strong relationship known as a “medium relationship”. This is because when the average person thinks of the root term “skiing” they do not typically think also of “sport”. Going from “ski” to “skiing” however, the average person would more likely associate or think “skiing” if given the term “ski”. The direction of the arrows in the bond strengths indicates the direction of association. For example, “A→B” in FIG. 3 means that given the concept A, how likely is it or

closely would an average person associate the concept B. Going the other direction between the same two concepts may produce a different bond strength.

[0034] A “weak relationship” is displayed between “ski” and “K2 ski” (when an average person thinks of “ski,” “K2 ski” does not closely come to mind). However, if an average person were to go from “K2 ski” to “ski” this might be construed as a strong relationship since they would naturally associate “ski” if given “K2 ski”.

[0035] FIG. 3 also shows semantic distances between concepts. “Ski” and “skiing” have only a distance of two between them while “skiing” and “sport” have a distance of five (seven minus two). The distance between “ski” and “sport” is seven. When traveling from parent to child or vice-versa, the distances can be simply added/subtracted but when changing the direction of travel, a penalty may be imposed upon the distance calculation. Take, for example, the distance between “ski” and “athletic equipment company”. Judging merely on a linear basis, the distance might be twelve. But since the path from “ski” to “athletic equipment” switches direction twice (it starts down to “K2 ski” and then across the lateral bond to “K2” and then up to “athletic equipment company”) a penalty or scaling factor can cause the distance between “ski” and “athletic equipment” to be much larger than just twelve especially given their lack of connectedness. As described above, penalties may be added when the direction of traversal is switched or when a lateral bond is crossed. Distances between concepts may be calculated and stored for future use. Semantic distances

can be determined in a variety of ways, such as, for example, as described in U. S. Patent No. 6,453,315, which is hereby incorporated by reference.

[0036] The concept processor 130 can update and expand the ontology to include new concepts or update relationships and other information for concepts already present. Concepts can be updated or added to reflect the changes in everyday culture or for a particular enterprise. For example, if a new ski manufacturer comes into existence called “Big Ski” this concept can be added to the ontology by a user. The user can define the relationships of other concepts to “Big Ski” and the strength of these relationships as well as other data that defines the concept.

[0037] Concepts can be presented to a user 112a on a client device 102a in a concept view, such as a concept window. Figure 4 illustrates an example of a concept window 400. The concept window 400 shows an example of a concept definition for “The United States”. The primary term is shown at 402 and is “The United States”. The primary term 402 also indicates the part of speech and provides a context for the concept (or refinition).

[0038] The concept window 400 shown further provides a concept details section 401 and a concept relationships section 412. The concept details section 401 can display various details of the concept. For example, at 404 the group of related terms associated with the concept is shown. The part of speech is shown at 406 and indicates that the concept shown is a noun. Refinition is showed at 408 and user data is shown at 410. The refinition box can be used to provide a context for the concept and the user data box can be used to further describe the concept.

[0039] The relationships for the concept are shown in the relationship section 412. The concept origin is shown at 414. The various relationships for the concept are shown in boxes 416, 418, and 420. In box 416, the parents of the concept are displayed and in box 420, children of the concept are displayed. Parent concepts and children concepts are concepts with a hierarchical relationship with the origin concept. In box 418, concepts having a lateral relationship with the origin concept are displayed

[0040] Various methods in accordance with the present invention may be carried out. For example, in one embodiment a user request to create a first concept in a network of interrelated concepts is received, a concept view, such as a concept creation window, is displayed, a primary term representing the first concept is received, at least one related term associated with the primary term is received, at least one relationship between the first concept and a second concept is received, a strength value associated with the relationship is received, and the first concept is added to the network. In one embodiment, the relationship can be hierarchical or lateral and can be selected from one of the following: kind of, has kind, part of, has part, member of, has member, substance of, has substance, product of, and has product.

[0041] In another embodiment, a request to edit a first concept in a network of interrelated concepts is received, the first concept is displayed, at least one relationship between the first concept and a second concept is received, a strength value associated with the relationship is received, an edited strength value for an

existing relationship between the first concept and another concept is received, and the first concept is updated in the network.

[0042] In one embodiment, information defining a part of speech of the first concept, information defining a frequency of the primary term, information defining a likelihood that the primary term and the related terms imply the first concept, information defining a breadth of the first concept, information that the first concept is offensive, user data, and context information, such as refinition, can also be received in order to create or edit the first concept.

[0043] FIG. 5 illustrates an exemplary method 500 that provides a method for creating concepts in a network of interconnected concepts or ontology. This exemplary method is provided by way of example, as there are a variety of ways to carry out methods according to the present invention. The method 500 shown in FIG. 5 can be executed or otherwise performed by any of various systems. The method 500 is described below as carried out by the system 100 shown in FIG. 1 by way of example, and various elements of the system 100 are referenced in explaining the example method of FIG. 5.

[0044] In 502, a request signal from a user to create a concept term is received by the semantic engine 120. The request signal can be created by a user interface application on client device 102a through input from the user 112a and passed to the server device 104 via the network 106.

[0045] In response to the signal, the concept processor 130 can cause, in 504, a concept view, such as a concept creation window, to be displayed on the client device

102a. The user 112a can be prompted for the entry of a variety of information in order to define the concept by the concept creation window. The information that defines the concept can include: a primary term used to identify the concept; a group of terms related to the primary term, such as synonyms and acronyms; related concepts and the relationship of the created concept to the related concepts; and the strength of these relationships. The information that defines the concept can also include: the part of speech of the concept; the frequency of the primary term; the likelihood that the primary term and the related terms imply the concept; the breadth or focus of the concept; the offensive nature of the concept; and the literal nature of the primary term. A concept definition can also include user data and a redefinition.

[0046] In defining a new concept, a user can input a primary term associated with the concept in the concept creation window. The primary term can be the term most frequently associated with the concept. Next, a user can input a group of terms related to the primary term that are also associated with the concept in the concept creation window. For example, the related terms can be synonyms or acronyms of the primary term.

[0047] A user 112a can also input concepts related to the created concept and provide relationship types for each related concept. The relationships can include: “kind of”; “has kind”; “part of”; “has part”; “member of”; “has member”; “substance of”; “has substance”; “product of”; and “has product”. Relationship types can also be specialized or customized, such as “causes symptom” and “symptom caused by” for a pharmaceutical or medical concept. The relationship type can also be “single

concept” for concepts unrelated to others. In one embodiment, the user 112a can define the related concept by typing it in a related concept box and then defining the relationship type by choosing a relationship type from a drop down list.

[0048] For every relationship type, the user 112a can assign a strength to the relationship. The strength can be representative of how strong the relationship is. In one embodiment, the strength represents the distance in semantic space the created concept is from the related concept.

[0049] The user 112a can also input other information to define the concept. For example, the part of speech of the concept can be provided. A frequency of the primary term can be provided to represent the frequency of the primary term in normal usage or in the source documents. A probability can be provided of the likelihood that the primary term and the related terms imply the concept. The user can specify the breadth or focus of the concept and indicate whether the concept is offensive. The user 112a can also specify whether to match terms associated with the concept literally or can define non-literal acronym and abbreviation rules. The user can also provide user data and context for the concept (or refinition).

[0050] After the user 112a has defined the concept, the user 112a can request through the client device 102a that the concept be added to the ontology. This request can cause the data input by the user 112a in the concept creation window to be sent from the client device 102a via the network 106 to the semantic engine 120. In 508, the concept processor 130 can create a new concept in the ontology from the data

received from the client device 102a and can add the concept to the ontology in the concept database 122.

[0051] FIG. 6 illustrates an exemplary method 600 that provides a method for editing concepts in an ontology. This exemplary method is provided by way of example, as there are a variety of ways to carry out methods according to the present invention. The method 600 shown in FIG. 6 can be executed or otherwise performed by any of various systems. The method 600 is described below as carried out by the system 100 shown in FIG. 1 by way of example, and various elements of the system 100 are referenced in explaining the example method of FIG. 6.

[0052] In 602, a request signal to edit a concept from a user 112a is received by the semantic engine 120. The request signal can be created by a user interface application on client device 102a through input from the user 112a. This request signal can include the concept to be edited and is passed to server device 104 via the network 106.

[0053] In response to the signal, the concept processor 130 can verify that the concept requested exists and can cause in 604 a concept editor window to be displayed on client device 102a. The concept editor window displays the concept definition that the user has requested to edit. A user can edit the various information used to define a concept as described above.

[0054] A user can edit the group of terms related to the primary term that are associated with the concept. For example, the user can add additional terms or delete related terms, such as synonyms or acronyms, of the primary term.

[0055] A user can also input additional concepts related to the edited concept, input relationship types for each related concept, and edit existing concepts and relationship types. In one embodiment, a user can define and/or edit the relationship type by choosing a relationship type from a drop down list. For the relationship types in the concept definition, the user can assign and edit a strength of the relationship. The user 112a can also input and edit other information in order to edit the concept as described above.

[0056] After the user 112a has edited the concept, the user 112a can request through the client device 102a that the edited concept be added to the ontology. This request can cause the data input by the user in the concept edit window to be sent from the client device 102a via the network 106 to the semantic engine 120. In 606, the edited concept is received by the semantic engine 120. In 608, the concept processor 130 can replace the concept in the ontology with the edited concept received from the user 112a or can add the edited data received from the user 112a to the concept in the ontology.

[0057] While the above description contains many specifics, these specifics should not be construed as limitations on the scope of the invention, but merely as exemplifications of the disclosed embodiments. Those skilled in the art will envision many other possible variations that are within the scope of the invention.